CLAIMS

We Claim:

1. A method comprising:

generating synthesized speech samples, using a synthesis filter, in response to an excitation signal;

determining a synthesis error between original speech samples and the synthesized speech samples; and

substantially reducing the synthesis error by computing both the excitation signal and filter parameters for the synthesis filter, wherein substantially reducing the synthesis error comprises applying a gradient descent algorithm to a polynomial representing the synthesis error over a series of iterations, including computing a gradient of the synthesis error in terms of gradient vectors of the synthesized speech samples by generating partial derivatives, using a recursive algorithm, for terms of a polynomial representing the synthesized speech samples over a series of iterations.

2. The method defined in Claim 1 wherein substantially reducing the synthesis error occurs in the root domain and the gradient descent algorithm is applied to roots of the polynomial.

3. The method defined in Claim 1 wherein substantially reducing the synthesis error comprises finding the roots of the polynomial representing the synthesis error.

4. The method defined in Claim 3 wherein finding the roots of a polynomial representing the synthesis error comprises converting linear predictive coding (LPC) coefficients to roots.

5. The method defined in Claim 1 wherein reducing the synthesis error between original speech samples and the synthesized speech samples occurs in the line spectrum pair (LSP) or the line spectrum frequency (LSF) domain, and generating partial derivatives, using a recursive algorithm, for terms of the polynomial representing the synthesized speech samples comprises computing partial derivatives with respect to line spectrum pairs (LSPs).

6. The method defined in Claim 5 wherein the LSPs comprise roots of a pair of polynomials based on line spectrum frequencies (LSFs).

7. The method defined in Claim 5 further comprising using gradient descent to optimize LSPs for the excitation signal to reduce an error between the original speech samples and the synthesized speech samples.

8. The method defined in Claim 5 wherein substantially reducing the synthesis error comprises finding the LSPs of the polynomial representing the synthesis error.

9. The method defined in Claim 1 further comprising adjusting a step-size used in the gradient descent algorithm at each iteration to ensure that a minimum of the synthesis error is not overshot.

10. The method defined in Claim 9 further comprising repeating use of a previous parameter vector from a previous iteration using a smaller step size if use of the current parameter vector from the current iteration causes the synthesis filter to become unstable or the synthesis error resulting from use of the current parameter vector is greater than the synthesis error resulting from use of the previous parameter vector.

11. The method defined in Claim 10 wherein adjusting the step-size continues until the synthesis filter regains stability or the synthesis error of the current iteration becomes smaller than the synthesis error of the previous iteration.

12. The method defined in Claim 1 further comprising generating an excitation function.

13. The method defined in Claim 12 wherein generating an excitation function comprises selecting the excitation function from a codebook of possible excitations.

14. An apparatus comprising:

an excitation unit to generate an excitation signal;

a synthesis filter to generate synthesized speech samples in response to the excitation signal;

a synthesis error generator to determine a synthesis error between original speech samples and the synthesized speech samples; and

a synthesis filter and excitation unit parameter generator to compute both the excitation signal and filter parameters for the synthesis filter in order to substantially reduce the synthesis error, wherein the synthesis filter and excitation unit parameter generator substantially reduces the synthesis error by applying a gradient descent algorithm to a polynomial representing the synthesis error over a series of iterations, including computing a gradient of the synthesis error in terms of gradient vectors of the synthesized speech samples by generating partial derivatives, using a recursive algorithm, for terms of a polynomial representing the synthesized speech samples over a series of iterations.

15. The apparatus defined in Claim 14 wherein the synthesis filter and excitation unit parameter generator substantially reduces the synthesis error in the root domain and the gradient descent algorithm is applied to roots of the polynomial.

16. The apparatus defined in Claim 14 wherein the synthesis filter and excitation unit parameter generator substantially reduces the synthesis error by finding the roots of the polynomial representing the synthesis error.

17. The apparatus defined in Claim 16 wherein the synthesis filter and excitation unit parameter generator finds the roots of the polynomial representing the synthesis error by converting linear predictive coding (LPC) coefficients to roots .

18. The apparatus defined in Claim 14 wherein the synthesis filter and excitation unit parameter generator reduces the synthesis error between original speech samples and the synthesized speech samples occurring in the line spectrum pair (LSP) domain, and generates partial derivatives, using a recursive algorithm, for terms of the polynomial representing the synthesized speech samples comprises computing partial derivatives with respect to line spectrum pairs (LSPs).

19. The apparatus defined in Claim 18 wherein the LSPs comprise roots of a pair of polynomials based on line spectrum frequencies (LSFs).

20. The apparatus defined in Claim 18 wherein the synthesis filter and excitation unit parameter generator uses gradient descent to optimize LSPs for the excitation signal to reduce an error between the original speech samples and the synthesized speech samples.

21. The apparatus defined in Claim 18 wherein the synthesis filter and excitation unit parameter generator substantially reduces the synthesis error by finding the LSPs of the polynomial representing the synthesis error.

22. The apparatus defined in Claim 14 wherein the synthesis filter and excitation unit parameter generator adjusts a step-size used in the gradient descent algorithm at each iteration to ensure that a minimum of the synthesis error is not overshot.

23. The apparatus defined in Claim 22 wherein the synthesis filter and excitation unit parameter generator repeats use of a previous parameter vector from a previous iteration using a smaller step size if use of the current parameter vector from the current iteration causes the synthesis filter to become unstable or the synthesis error resulting from use of the current parameter vector is greater than the synthesis error resulting from use of the previous parameter vector.

24. The apparatus defined in Claim 14 wherein the excitation unit generates an excitation signal by selecting an excitation function from a codebook of possible excitations.

25. An article of manufacture having one or more recordable media storing instructions which, when executed by a system, cause the system to:

generate synthesized speech samples, using a synthesis filter, in response to an excitation signal;

determine a synthesis error between original speech samples and the synthesized speech samples; and

substantially reduce the synthesis error to compute both the excitation signal and filter parameters for the synthesis filter, wherein the instructions to substantially reduce the synthesis error comprise instruction which, when executed by a system, cause the system to apply a

gradient descent algorithm to a polynomial representing the synthesis error over a series of iteration, including computing a gradient of the synthesis error in terms of gradient vectors of the synthesized speech samples by generating partial derivatives, using a recursive algorithm, for terms of a polynomial representing the synthesized speech samples over a series of iterations.

26. The article of manufacture defined in Claim 25 wherein the substantial reduction in the synthesis error occurs in the root domain and the gradient descent algorithm is applied to roots of the polynomial.

27. The article of manufacture defined in Claim 25 wherein the instructions to substantially reduce the synthesis error comprise instructions which, when executed by a system, cause the system to find the roots of the polynomial representing the synthesis error by converting linear predictive coding (LPC) coefficients to roots.

28. The article of manufacture defined in Claim 27 wherein the instructions to generate partial derivatives, using a recursive algorithm, for terms of the polynomial representing the synthesized speech samples over a series of iterations comprises instructions which, when executed by a system, cause the system to generate the partial derivatives for each root of a polynomial representing the synthesized speech samples during an iteration in the series of iterations.

29. The article of manufacture defined in Claim 25 wherein reduction of the synthesis error between original speech samples and the synthesized speech samples occurs in the line

spectrum pair (LSP) domain, and instructions to generate partial derivatives, using a recursive algorithm, for terms of the polynomial representing the synthesized speech samples comprise instructions which, when executed by a system, cause the system to compute partial derivatives with respect to line spectrum pairs (LSPs).

30. The article of manufacture defined in Claim 29 wherein the LSPs comprise roots of a pair of polynomials based on line spectrum frequencies (LSFs).

31. The article of manufacture defined in Claim 30 wherein instructions to substantially reduce the synthesis error comprises instructions which, when executed by a system, cause the system to find the LSPs of the polynomial representing the synthesis error.

32. The article of manufacture defined in Claim 31 wherein instructions to compute the gradient vector comprise instructions which, when executed by a system, cause the system to compute the gradient vector of the synthesis error relative to LSP vectors.

33. The article of manufacture defined in Claim 32 wherein instructions to generate partial derivatives, using a recursive algorithm, for terms of the polynomial representing the synthesized speech samples over a series of iterations comprise instructions which, when executed by a system, cause the system to generate the partial derivatives for each LSP of a polynomial representing the synthesized speech samples during an iteration in the series of iterations.